Linear Filtering

Computer Vision (CSCI 4220U)

Faisal Z. Qureshi
http://vclab.science.ontariotechu.ca
Ontario Tech University

Readings

- Szeliski 2nd Edition, Section 3.2
- Checkout course notes at http://csundergrad.science.uoit.ca/courses/cv-notes/notebooks/03-linear-filtering.html

Linear Filtering in 1D

$\mathbf{f} =$

1	3	4	1	10	3	0	1

$\mathbf{h} =$

1	0	-1

$(w = 1)$

(4)(1) + (1)(0) + (10)(-1)

	-3	2	-6				

Cross-correlation

Linear Filtering in 1D

$\mathbf{f} =$

1	3	4	1	10	3	0	1

$\mathbf{h} =$

1	0	-1

$(w = 1)$

Cross-Correlation

$$\mathrm{CC}(i) = \sum_{k \in [-w, w]} \mathbf{f}(i+k)\mathbf{h}(k)$$

	-3	2	-6	-2	10	2	

Linear Filtering in 1D

$\mathbf{f} =$

1	3	4	1	10	3	0	1

$\mathbf{h} =$

1	0	-1

$(w = 1)$

Cross-Correlation

$$\mathrm{CC}(i) = \sum_{k\in[-w,w]} \mathbf{f}(i+k)\mathbf{h}(k)$$

	-3	2	-6	-2	10	2	

Convolution

$$(f * h)_i = \sum_{k\in[-w,w]} \mathbf{f}(i-k)\mathbf{h}(k)$$

	3	-2	6	2	-10	-2	

Filter $\mathbf{h}$ is flipped

Linear Filtering in 1D

$\mathbf{f} =$

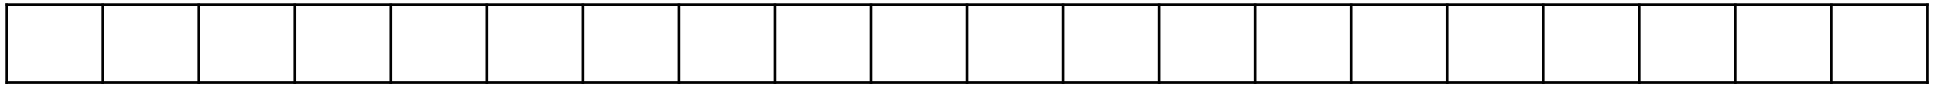

$\mathbf{h} =$

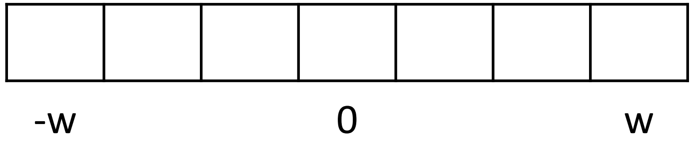

Total width of the filter/kernel is $(2w + 1)$

Cross-Correlation

$$\mathrm{CC}(i) = \sum_{k \in [-w,w]} \mathbf{f}(i+k)\mathbf{h}(k)$$

Convolution

$$(f * h)_i = \sum_{k \in [-w,w]} \mathbf{f}(i-k)\mathbf{h}(k)$$

Filter $\mathbf{h}$ is flipped

Same when the filter/kernel $\mathbf{h}$ is *symmetric*

Linear Filtering in 2D

1	1	1	2	1	1	1	4
1	1	2	2	2	1	1	4
1	2	2	2	2	2	1	4
1	3	3	0	0	3	1	4
2	1	1	0	1	1	1	2
1	64	3	22	1	32	1	7
8	5	7	4	2	2	8	9
8	8	9	8	0	0	0	0

Image

1	1	1
1	1	1
1	1	1

Filter

Linear Filtering in 2D

1	1	1	2	1	1	1	4
1	1	2	2	2	1	1	4
1	2	2	2	2	−2	1	4
1	3	3	0	0	3	1	4
2	1	1	0	1	1	1	2
1	64	3	22	1	32	1	7
8	5	7	4	2	2	8	9
8	8	9	8	0	0	0	0

11

Image

1	1	1
1	1	1
1	1	1

Filter

Summing kernel

Linear Filtering in 2D

1	1	1	2	1	1	1	4
1	1	2	2	2	1	1	4
1	2	2	2	2	−2	1	4
1	3	3	0	0	3	1	4
2	1	1	0	1	1	1	2
1	64	3	22	1	32	1	7
8	5	7	4	2	2	8	9
8	8	9	8	0	0	0	0

1.22

Image

$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

Filter

9 by 9 averaging kernel

Linear Filtering in 2D

Cross-Correlation

$$\mathrm{CC}(i,j) = \sum_{k\in[-w,w],l\in[-h,h]} \mathbf{f}(i+k, j+l)\mathbf{h}(k,l)$$

Convolution

$$(f * h)_{ij} = \sum_{k\in[-w,w],l\in[-h,h]} \mathbf{f}(i-k, j-l)\mathbf{h}(k,l)$$

Filter $\mathbf{h}$ is flipped, both horizontally and vertically

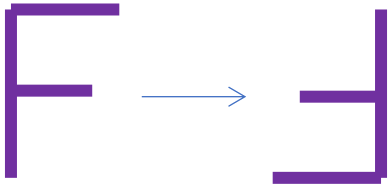

Convolution and cross-correlation is the same for symmetric kernels

$\mathbf{h} =$

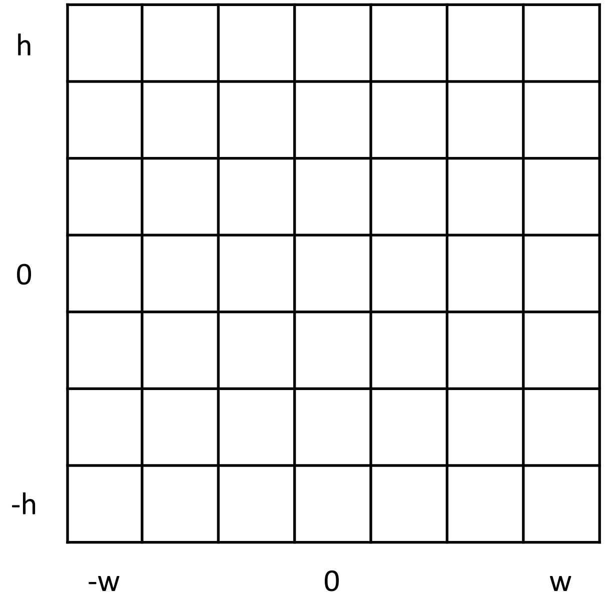

Cross-Correlation & Convolution

- Linear shift-invariant operations
- Superposition principle

$$h \circ (f_0 + f_1) = h \circ f_0 + h \circ f_1$$

- Shift-invariance principle

$$g(i,j) = f(i+k, j+l) \Leftrightarrow (h \circ g)(i,j) = (h \circ f)(i+k, j+l)$$

Padding

- Zero
- Constant
- Clamp or replicate
- Cyclic wrap
- Mirror

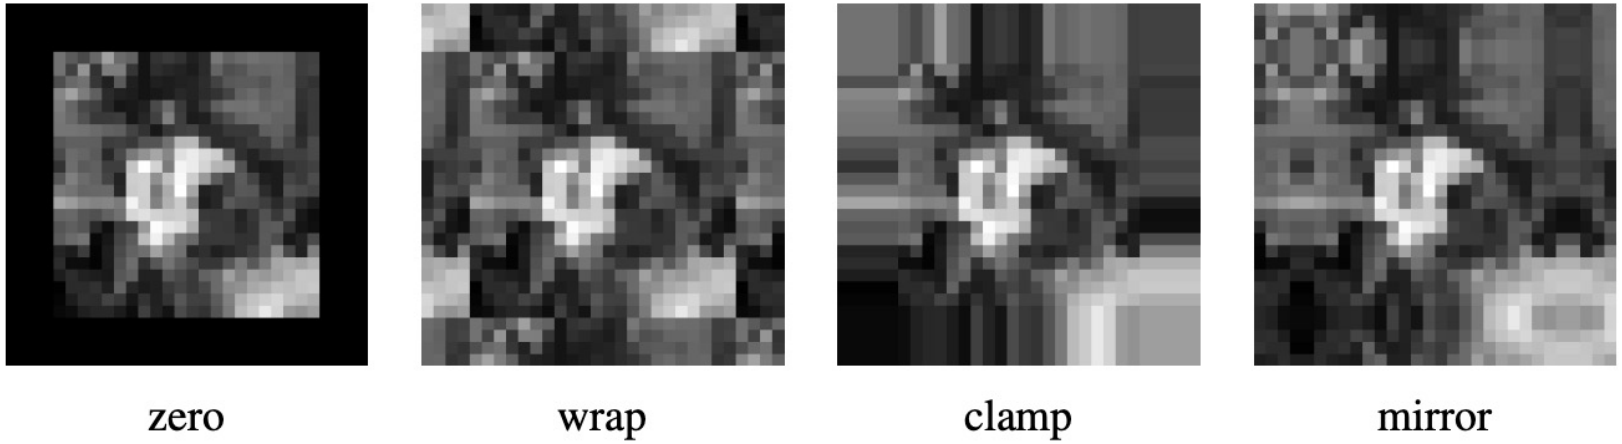

Separable Filters

- An n-dimensional filter that can be expressed as an outer-product of n 1-dimensional filters

$$\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$$

$\in R^{3\times 3}$ $\in R^{3\times 1}$ $\in R^{1\times 3}$

Separable Filters

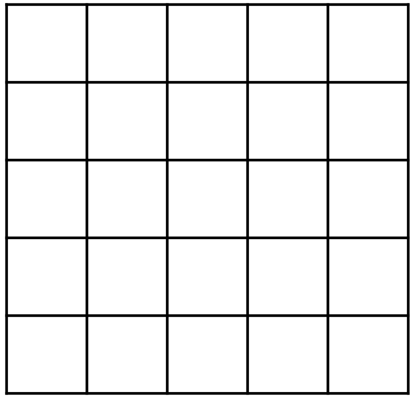

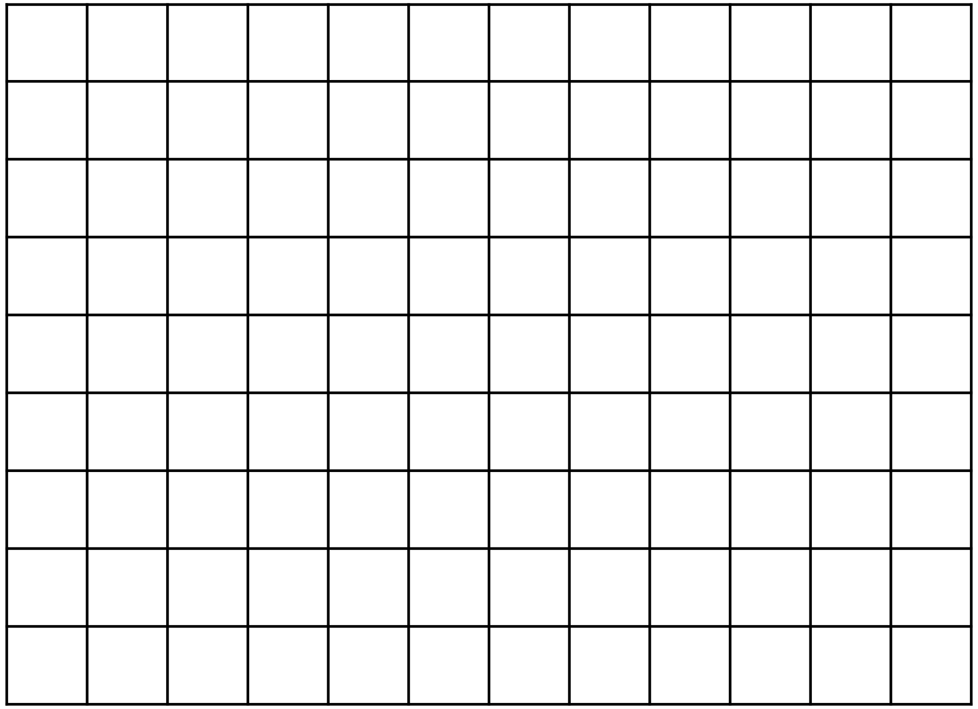

Separable Filters

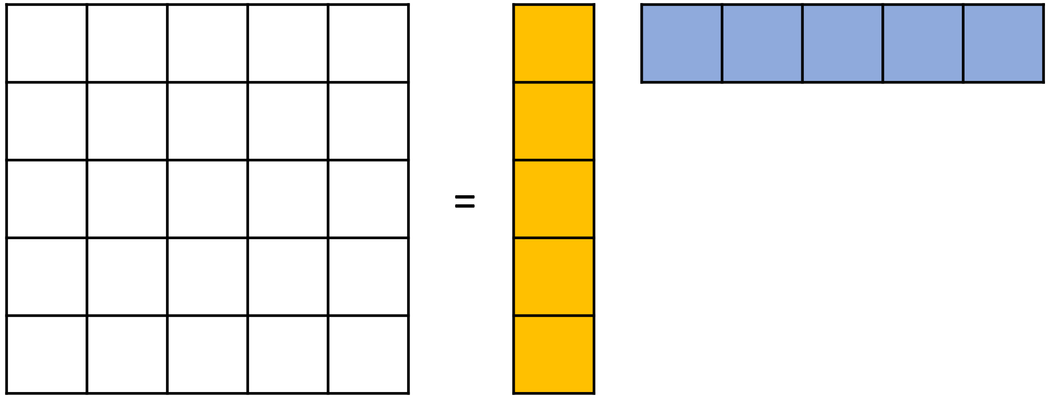

Convolution with Separable Filters in 2D

- [Step 1] Perform row-wise convolution with horizontal filter
- [Step 2] Perform column-wise convolution the results obtained in step 1 with vertical filter

Separable Filters

Step 1

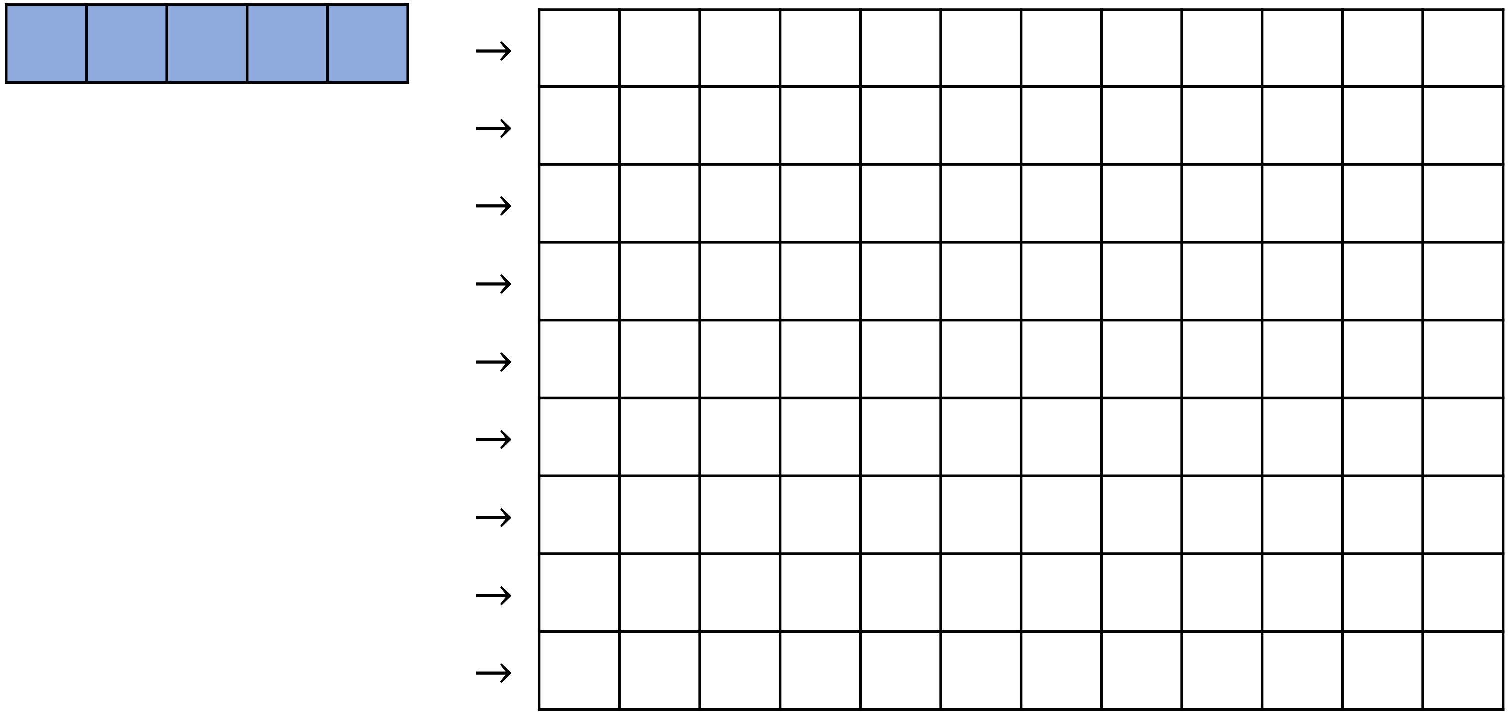

Separable Filters

Step 2

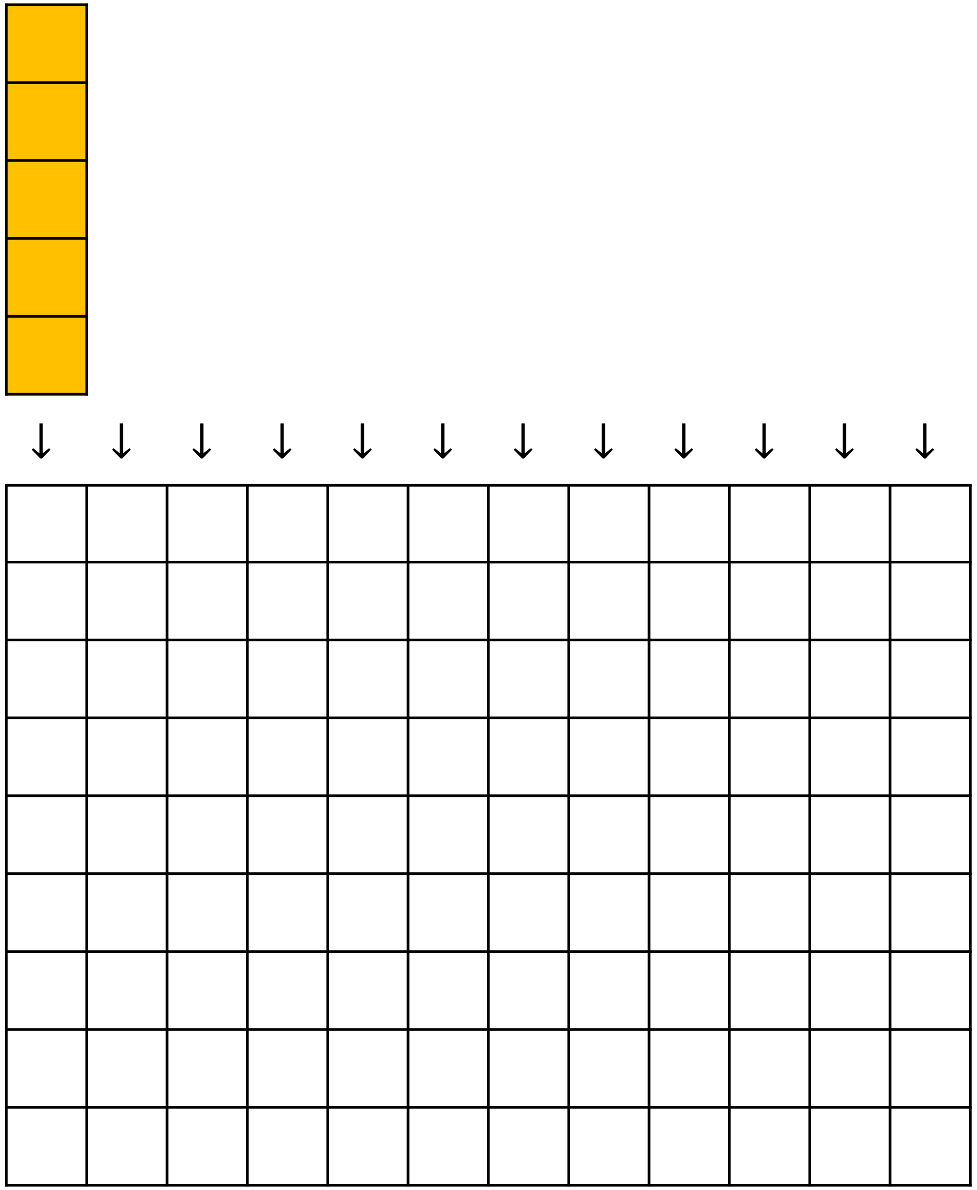

Separable Filters Example

Compute

1	2	1
2	4	2
1	2	1

$*$ (denotes convolution)

2	3	3
3	5	5
4	4	6

Recall that

1	2	1
2	4	2
1	2	1

=

1
2
1

1	2	1

Separable Filters Example

Compute

1	2	1
2	4	2
1	2	1

$*$

2	3	3
3	5	5
4	4	6

$= 65$

Recall that

1	2	1
2	4	2
1	2	1

$=$

1
2
1

1	2	1

Step 1

1	2	1

$*$

2	3	3
3	5	5
4	4	6

$=$

11
18
11

Step 2

1
2
1

$*$

11
18
11

$= 65$

Computational considerations

- Separable filters

$$O(w_k wh) + O(h_k wh)$$

- Non-separable filters

$$O(w_k h_k wh)$$

w_k, h_k, w, and h denote kernel width and height and image width and height, respectively.

Separable filters

- Use Singular Value Decomposition (SVD) to determine if a filter is separable
- If only one singular value is non-zero then the filter is separable

$$\mathbf{F} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_i \sigma_i \mathbf{u}_i \mathbf{v}_i^T$$

  with $\Sigma = \mathrm{diag}(\sigma_i)$

- Vertical and horizontal filters are $\sqrt{\sigma_1}\mathbf{u}_1$ and $\sqrt{\sigma_1}\mathbf{v}_1^T$, respectively

- Whenever possible use separable filters

Gaussian in 1D

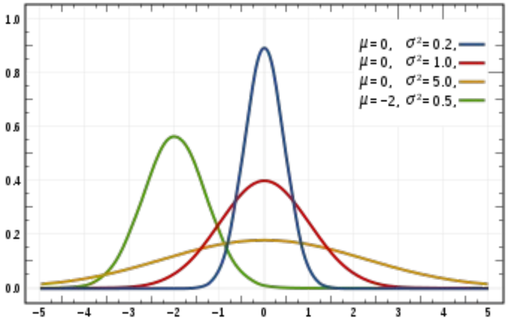

$$G(x; \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

Here μ and σ refer to the mean and standard deviation of this Gaussian.

Aside: Mean and Standard Deviation

Given n data points $\{x_1, x_2, \cdots, x_n\}$

$$\mu = \mathrm{E}[x] = \frac{1}{n} \sum_{i=1}^{n} x_i$$

$$\sigma^2 = \mathrm{E}[(x-\mu)^2] = \frac{1}{n} \sum_{i=1}^{n} (x_i - \mu)^2$$

Multivariate Guassian (in k-dimensions)

$$G(\boldsymbol{x};\boldsymbol{\mu},\Sigma)=\frac{1}{\sqrt{(2\pi)^k\det(\Sigma)}}exp\left(-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu})^T\Sigma^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right)$$

where

$\boldsymbol{x} \in R^k$

$\boldsymbol{\mu} \in R^k$

$\Sigma \in R^{k \times k}$

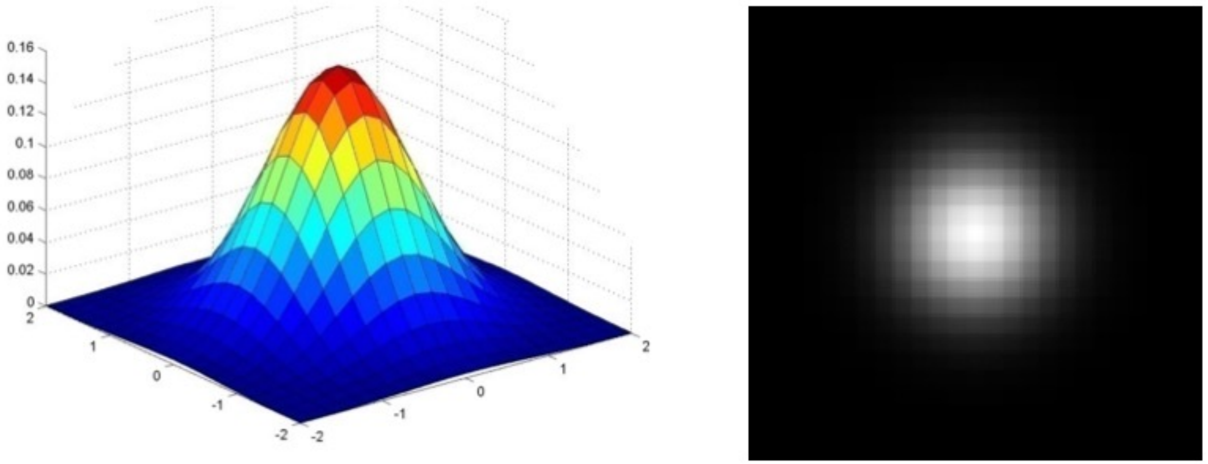

Gaussian in 2D

Gaussian Blurring

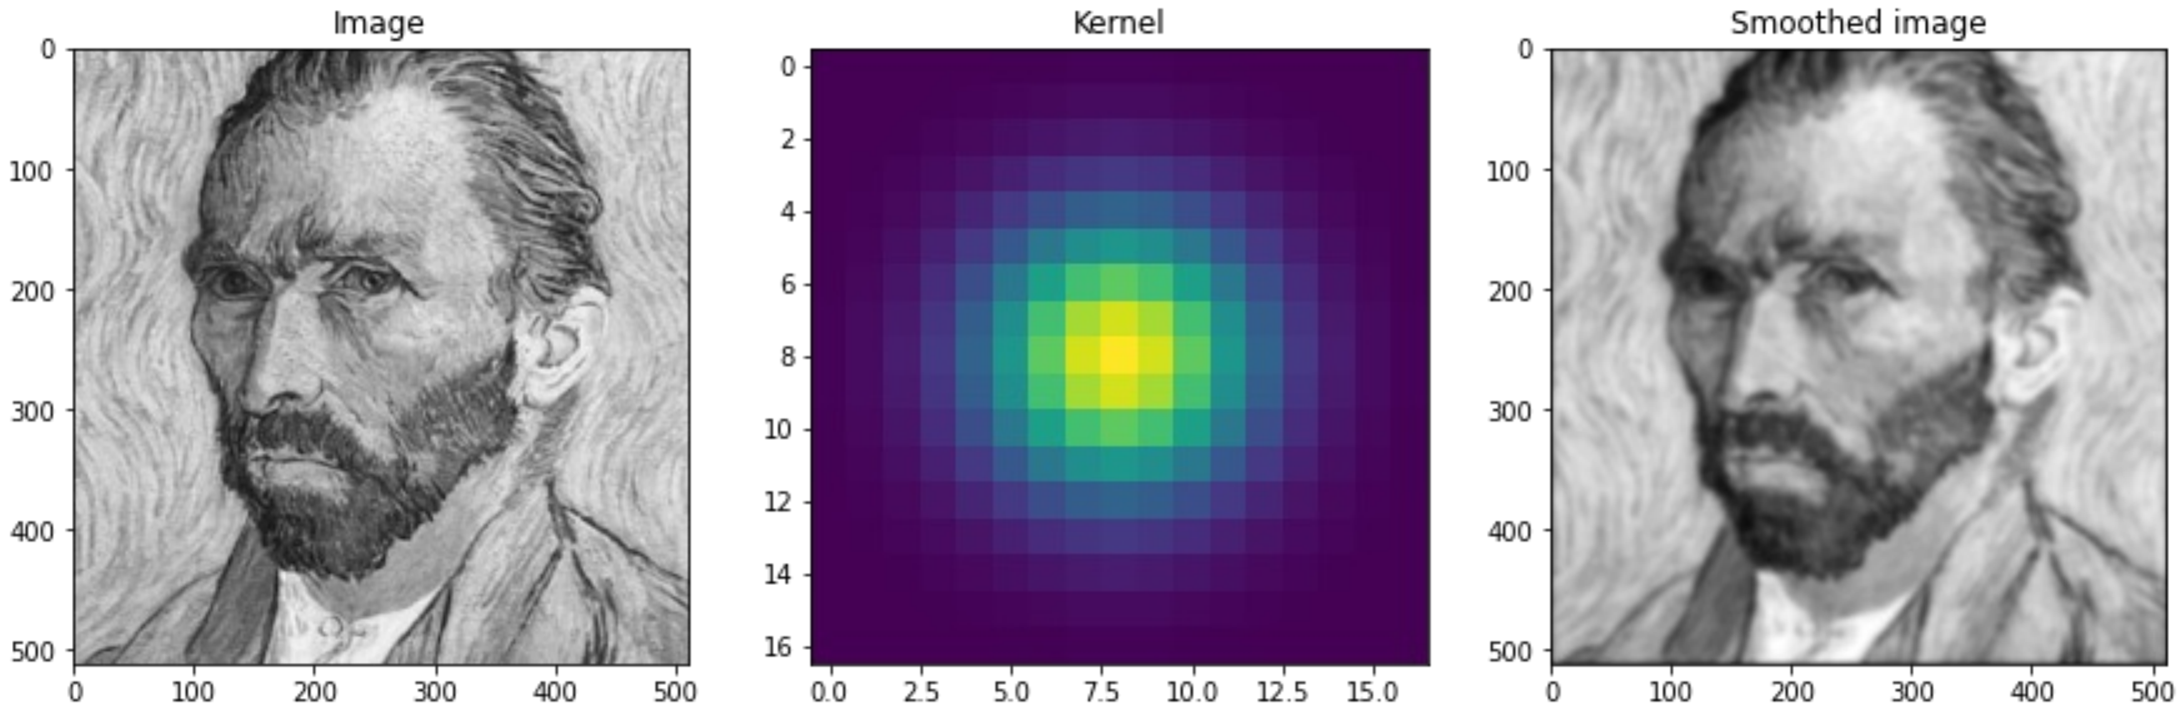

Gaussian Filter

- We often approximate 2D Gaussian filter as follows

$$G(k, l; 0, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{k^2+l^2}{2\sigma^2}}$$

- Example 3x3 Gaussian filter: $\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$

- Example 1x5 Gaussian filter is: $\frac{1}{16}[\,1 \quad 4 \quad 6 \quad 4 \quad 1\,]$

Gaussian Filter

- Gaussian functions have infinite support, but discrete Gaussian kernels are finite

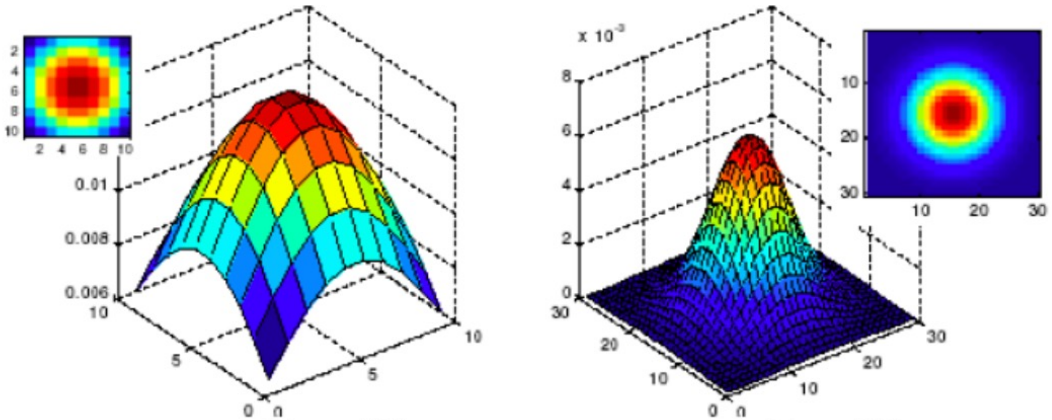

σ = 5 with 10 x 10 kernel

σ = 5 with 30 x 30 kernel

Gaussian Filter

- Variance controls how broad or peaky the filter is

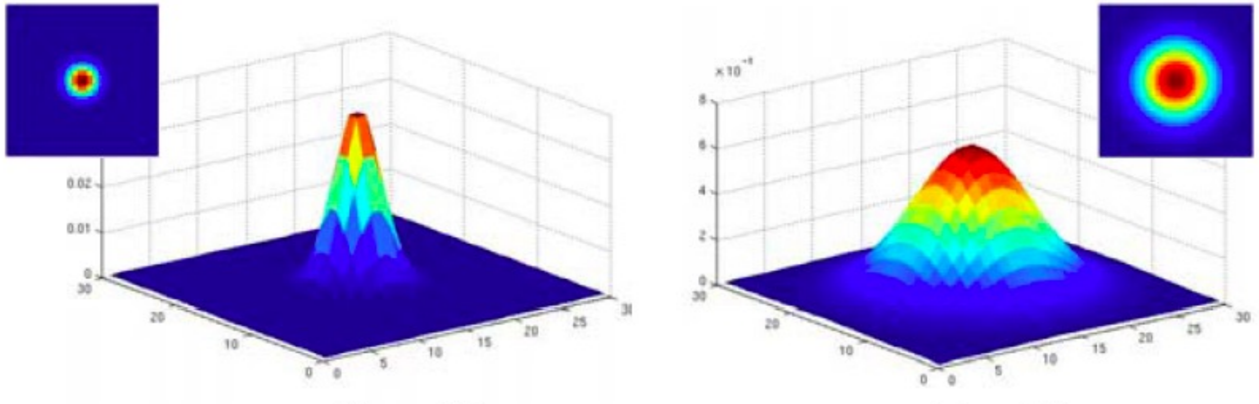

σ = 2 with 30 x 30 kernel

σ = 5 with 30 x 30 kernel

Gaussian Blurring

- Removes high-frequency components from the image
  - Blurs the image
  - Acts as a low-pass filter

Gaussian Blurring vs. Average (Box) Filtering

Gaussian Kernel

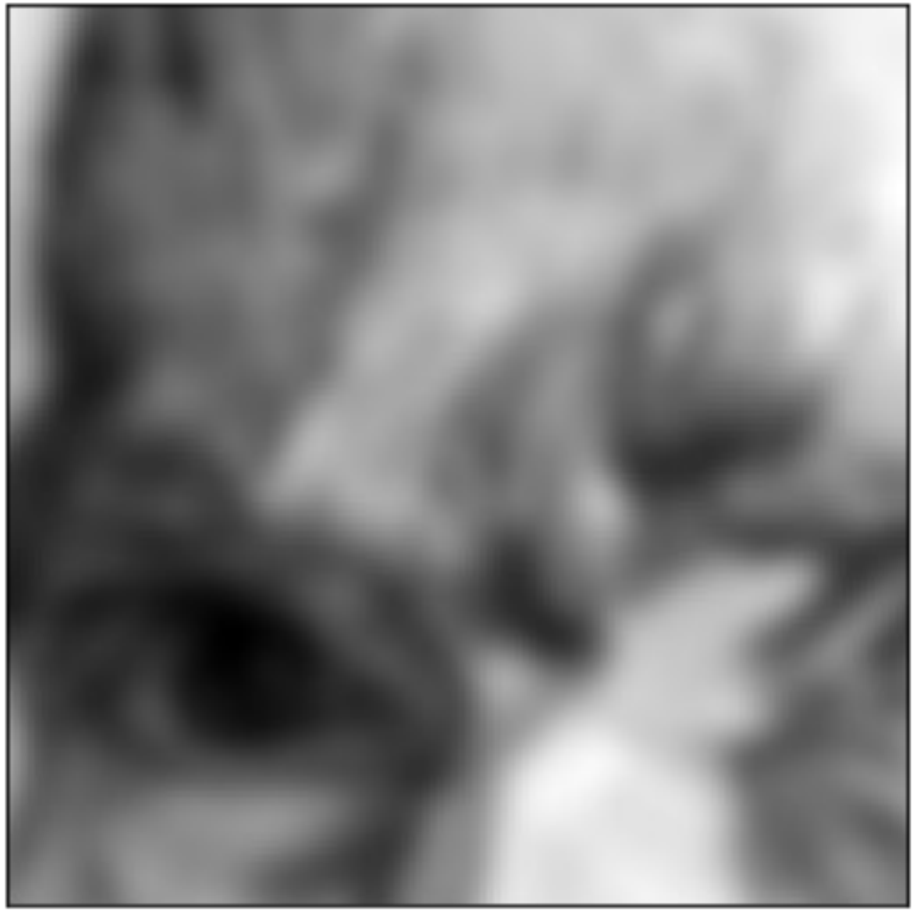

Averaging (Box) Kernel

Gaussian Blurring vs. Average (Box) Filtering

Gaussian Kernel

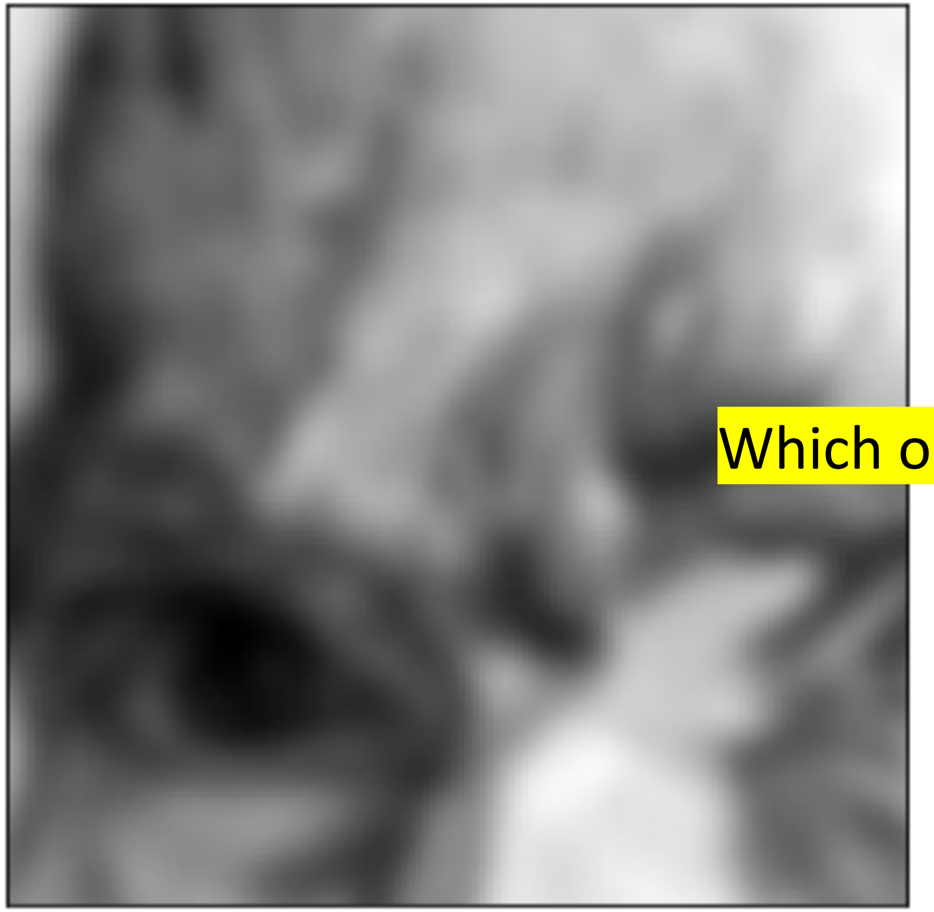

Averaging (Box) Kernel

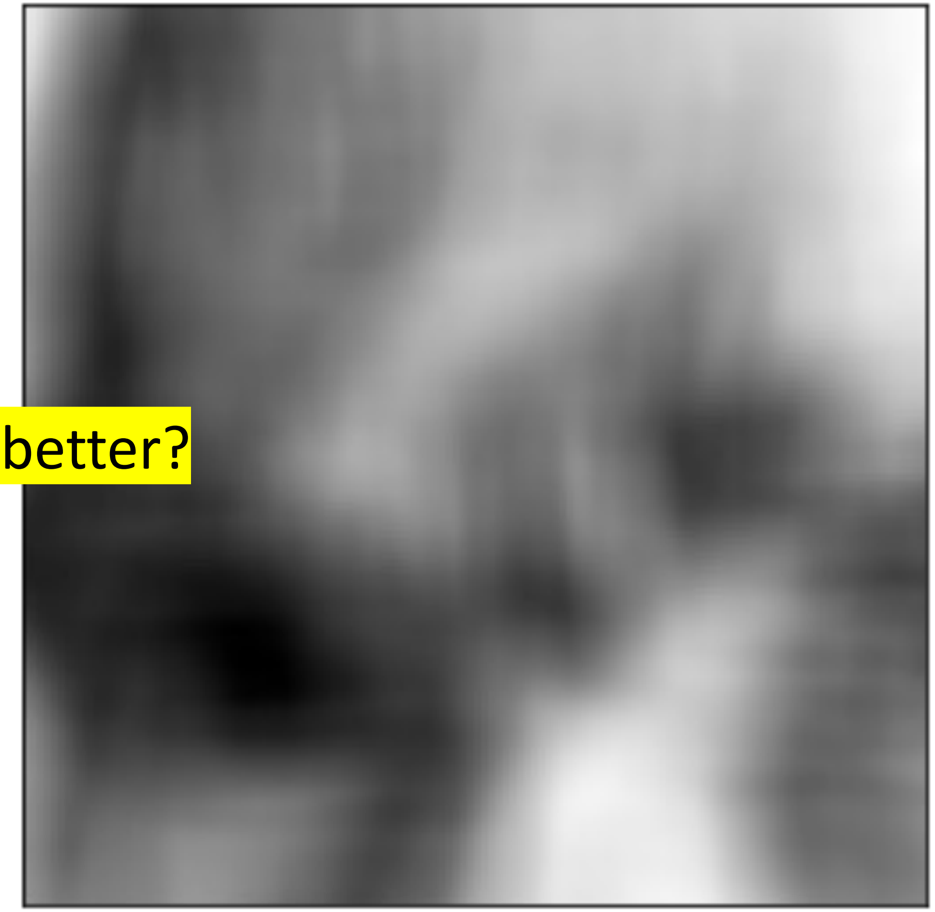

Which one is better?

Gaussian filter is separable

$$
\begin{aligned}
G_\sigma(x,y) &= \frac{1}{2\pi\sigma^2}\exp\left(-\frac{x^2+y^2}{2\sigma^2}\right) \\
&= \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{x^2}{2\sigma^2}\right)\cdot\frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{y^2}{2\sigma^2}\right) \\
&= g_\sigma(x)\cdot g_\sigma(y)
\end{aligned}
$$

FIGURE 24-7
Separation of the Gaussian. The Gaussian is the only PSF that is circularly symmetric *and* separable. This makes it a common filter kernel in image processing.

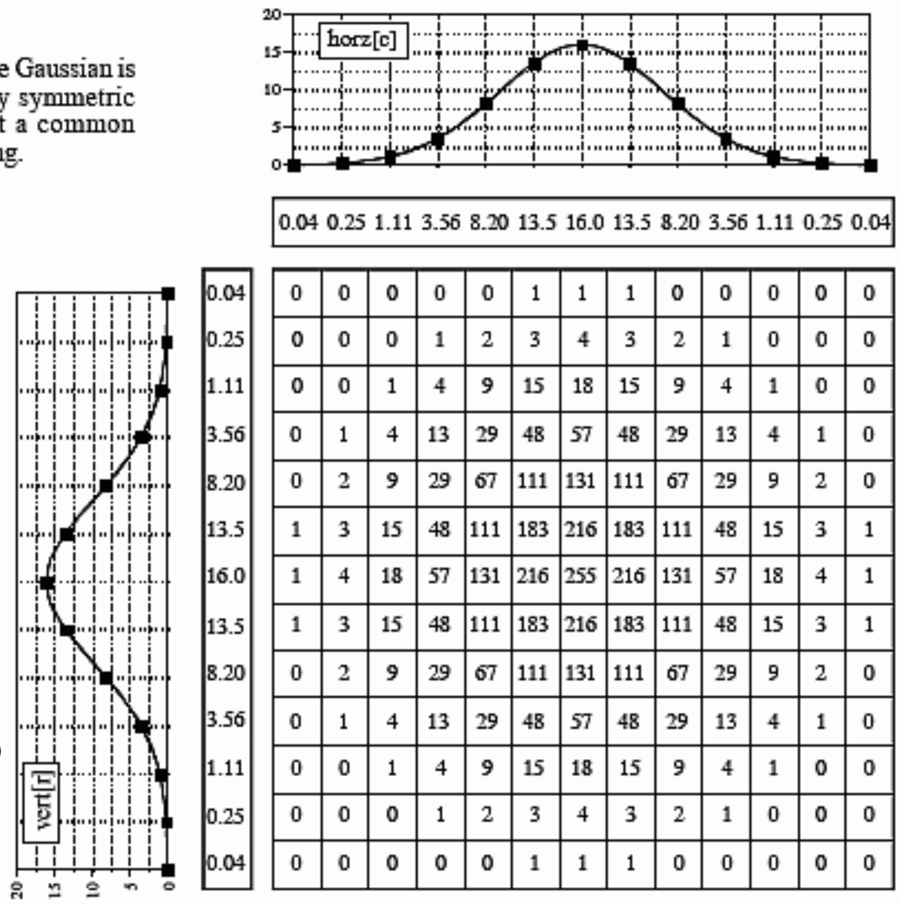

vert[r] \ horz[c]	0.04	0.25	1.11	3.56	8.20	13.5	16.0	13.5	8.20	3.56	1.11	0.25	0.04
0.04	0	0	0	0	0	1	1	1	0	0	0	0	0
0.25	0	0	0	1	2	3	4	3	2	1	0	0	0
1.11	0	0	1	4	9	15	18	15	9	4	1	0	0
3.56	0	1	4	13	29	48	57	48	29	13	4	1	0
8.20	0	2	9	29	67	111	131	111	67	29	9	2	0
13.5	1	3	15	48	111	183	216	183	111	48	15	3	1
16.0	1	4	18	57	131	216	255	216	131	57	18	4	1
13.5	1	3	15	48	111	183	216	183	111	48	15	3	1
8.20	0	2	9	29	67	111	131	111	67	29	9	2	0
3.56	0	1	4	13	29	48	57	48	29	13	4	1	0
1.11	0	0	1	4	9	15	18	15	9	4	1	0	0
0.25	0	0	0	1	2	3	4	3	2	1	0	0	0
0.04	0	0	0	0	0	1	1	1	0	0	0	0	0

The Scientist and Engineer's Guide to Digital Signal Processing
By Steven W. Smith, Ph.D.

Gaussian Filter

- How would you construct a 5x5 Gaussian filter given a the following 1x5 Gaussian filter?

$$\frac{1}{16}[\,1 \quad 4 \quad 6 \quad 4 \quad 1\,]$$

- Hint: employ the fact that the Gaussian filter is separable.

Examples of Separable Filters

(a) box, $K = 5$:

$$\frac{1}{K^2}\begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 \\ \vdots & \vdots & 1 & \vdots \\ 1 & 1 & \cdots & 1 \end{bmatrix} \qquad \frac{1}{K}\begin{bmatrix} 1 & 1 & \cdots & 1 \end{bmatrix}$$

(b) bilinear:

$$\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} \qquad \frac{1}{4}\begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$$

(c) "Gaussian":

$$\frac{1}{256}\begin{bmatrix} 1 & 4 & 6 & 4 & 1 \\ 4 & 16 & 24 & 16 & 4 \\ 6 & 24 & 36 & 24 & 6 \\ 4 & 16 & 24 & 16 & 4 \\ 1 & 4 & 6 & 4 & 1 \end{bmatrix} \qquad \frac{1}{16}\begin{bmatrix} 1 & 4 & 6 & 4 & 1 \end{bmatrix}$$

(d) Sobel:

$$\frac{1}{8}\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} \qquad \frac{1}{2}\begin{bmatrix} -1 & 0 & 1 \end{bmatrix}$$

(e) corner:

$$\frac{1}{4}\begin{bmatrix} 1 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 1 \end{bmatrix} \qquad \frac{1}{2}\begin{bmatrix} 1 & -2 & 1 \end{bmatrix}$$

(a) box, $K = 5$ (b) bilinear (c) "Gaussian" (d) Sobel (e) corner

Figure 3.14 *Separable linear filters: For each image (a)–(e), we show the 2D filter kernel (top), the corresponding horizontal 1D kernel (middle), and the filtered image (bottom). The filtered Sobel and corner images are signed, scaled up by 2× and 4×, respectively, and added to a gray offset before display.*

Self-Study

- Band-pass and steerable filters
- Summed area images (integral images)

Summary

- Linear filtering
- Cross-correlation
- Convolution
- Separable filters
- Gaussian filter